

An Interview with

DAVID WALDEN

OH 181

Conducted by Judy O'Neill

on

6 February 1990

Cambridge, MA

Charles Babbage Institute
The Center for the History of Information Processing
University of Minnesota, Minneapolis

Copyright, Charles Babbage Institute

David Walden Interview
6 February 1990

Abstract

Following a brief overview of his background, Walden traces his involvement with the ARPANET, from discussions before the official DARPA request was issued to his later management of the project at Bolt Beranek and Newman (BBN). He discusses the people involved in the ARPANET work at BBN and how they were influenced by their previous work developing real-time computer systems at Lincoln Laboratory. Walden describes the working environment of the group at BBN and their relationships with the IPT Office, other DARPA contractors, and the larger community. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

David Walden Interview
6 February 1990

Tape Index

Tape 1/Side 1.....	1
Side 2.....	14
Tape 2/Side 1.....	26

WALDEN: This is a picture of the ARPANET team at BBN. There are a couple of people missing. This is Bob Kahn, who went down to DARPA later to run the project. This is Frank Heart who was the project leader. This is me. Will Crowther, who was the lead software person. Barker who was one of the hardware people. Severo is right here; you can barely see him. And there are a couple of people who are not in the picture - Hawley Rising and ~~Bernie Cosell~~. But I can get you a print of that picture. I should make a list of things I'm going to do. Let me give you this. This is a paper that Alex and I wrote on the history of the ARPANET project. And this is going to be published in the *Encyclopedia of Telecommunications*. This summarizes the series of discussions in the ARPA community before the RFP came out as they derived what should go into the RFP. What was thought, was the packet switching program going to be in the host computers, was it going to be in its own computer, and so on. And some of the things that became quite clear very early on, in answer to your question of how have things changed a little bit, is for instance it became quite clear very early on, I am not sure if it was part of the RFP or not, but that there had to be multiple hosts on a packet switch. Originally, I think, there was going to be one or something. So that's the kind of thing. I can't remember in detail what it was.

WALDEN: Incidentally, have you got a copy of Larry Roberts' paper, which is in Adele Goldberg's book on Personal Workstations? Because that is another view. I think that view and this view are quite complementary. They are certainly not in disagreement and they cover the same material but not exactly. So the sum of them is quite a good summary. Did I give you or your colleague a copy of the original version of the ARPANET Completion Report, the big thick one?

O'NEILL: Yes.

WALDEN: In there, there is a discussion that I took down or that Bob Taylor sent to me when we wrote the ARPANET Completion Report. That is his view of how it started. You might look up in there (I'm not sure if it is presented as a quote or if it is a paraphrase or what), but what Bob Taylor said about the start of the

ARPANET in that three-volume book is worth looking into for another person's point of view. He talks about getting Roberts to come from Lincoln Labs and all of that.

O'NEILL: He was also interviewed for this project and discussed some of these things.

WALDEN: By the way, I think this article is a very good summary (or the three-volume one intended for the general public as opposed to the specialized public). The article carries on beyond the Completion Report into the last decade, what's happened with the ARPANET, how it is being dismantled, and so on. So it really is the history of the ARPANET and then how it affected the world, how it is now being dismantled, and how a number of the protocols and the internetwork, and so on, are going to live on.

O'NEILL: OK. Those are some of the things I wanted to discuss.

WALDEN: Alex McKenzie, by the way, is probably the best person to talk to about the ARPANET and its impact in the last ten years. He is the co-author of this paper and was one of the co-authors of the original ARPANET Completion Report.

O'NEILL: OK.

WALDEN: Can I interrupt and ask another question?

O'NEILL: Sure.

WALDEN: This can be informal.

O'NEILL: Yes, it is informal.

WALDEN: Is the thing you are working on here primarily the ARPANET or primarily ARPA technology?

O'NEILL: The thing I am personally working on is the ARPANET. The project we are doing involves other parts of it as well. But I am just interested right now in the ARPANET.

WALDEN: And your colleague has all the material on the other ARPA accomplishments that I gave him a year or so ago, long lists of what my view is on what they have done?

O'NEILL: Yes, a list of some of the programs. OK. Back to the response to the - is it RFP or is it RFQ?

WALDEN: It is called a RFQ, it is the same thing.

O'NEILL: Were there major differences between what you responded with and what other companies responded with?

WALDEN: I have no way of knowing, I assume so. I think that someone on the evaluation committee, like Roberts (I don't know who else was on the evaluation committee), would be in a better position to say that, but I suspect, for instance, that IBM bid big IBM machines (this is total speculation). I suspect that Adams Associates, which is a company which existed then, probably bid DEC machines, because they had a close association with DEC; Raytheon might have bid MILSPEC machines; I have no idea what anybody else did.

O'NEILL: In terms of your personal experience when you were with Lincoln Labs, did you have communications experience there, in the Space Communications Division?

WALDEN: I did. I think that my role in the ARPANET program was that I was the first computer programmer added to the design group. There was Frank Heart, who was the manager of the design group.

There was Bob Kahn, who was the leading theoretician, let's call it; I mean that in a fairly pragmatic sense of the word theoretician. He was the one who had been going to the design meetings in the community for the previous several years or however long it went on. His Ph.D. was in communications theory and so on and he had thought a lot about the issues and the architecture of the network. I read somewhere yesterday, some national magazine was saying something along the line that he basically had the architecture in his head. Smart guy. Severo who was the leading hardware guy was quite an experienced systems builder. I was programmer with four years' programming experience. I had participated in one substantial system at Lincoln Labs called Lincoln Experimental Terminal working with Frank Heart and Will Crowther. Will Crowther was my supervisor on that project. Frank was the assistant department manager for the department I worked in. So I would say that I was not a particularly experienced person. Because I was the first programmer on the project, I did a lot of the original thinking about the timing of the program, how long it would take to do things, how many instructions you would have to write. This was before we got the RFP. And I think that probably helped. I did some of the original block diagrams and that sort of thing. As it grew near the time when the RFP was about to come out (I am not exactly sure of the timing - Frank or Severo might remember it more) - I'm not sure exactly how to put it, but there was no doubt that we needed a more senior software person on the project than I was, and we had all been thinking of an excuse of why Willy should come from Lincoln. We were several of us here, he had worked with us there. Severo had been at Lincoln before he went to St. Louis to Washington University and was a friend of Willy's. Willy was an outstanding technical person, and so at some point along the line the overture was made to Willy. We were pretty sure he had not been completely happy since Frank left Lincoln. [We were in a group at Lincoln Labs that was a digital computing group inside an analog communications division. And Frank was the assistant department leader for that digital communications group. He worked at the digital interface. He had hardware engineers who worked at interface into the analog equipment and then the people who wrote the computer programs and did some of thinking of how the architecture should be with the other analog people. These were systems to point radar antennas and so on, so there were analog pieces, servo mechanisms and wave guides, and all that kind of stuff. When Frank left to come here, he was attracted here by Danny Bobrow, who is now at Xerox PARC. Danny was a division leader

here. And in fact Danny was one of the people that kibitzed the original ARPANET design effort and helped here when we were writing the proposal. Danny and Joe Marcowitz, who is now, I believe, working for the government in Washington, were two people who kibitzed the proposal or reviewed it. Bob Kahn reported in to Danny's division. He was sent from Danny's division to Frank Heart's division when it was decided to pull the project together in Frank's division.] Eventually Bob transferred into Frank's division. Anyway, when Frank left Lincoln Labs, this digitally-oriented group began to fall apart. People began drifting away. In the long run several others came here. I followed Frank first because I was the youngest and had the least invested. When this project came up we suggested that Willy should come.

So then Willy joined the project. He was obviously the senior software person on the project. In intellect, of the same calibre (different but of the same calibre) as Bob Kahn - really a brilliant guy; and so he and Bob would argue about algorithms and this and that. And Willy and I, as we had done at Lincoln Labs when I was his assistant just out of college, basically became the software development team. He did some things, I did some things. Certainly I kibitzed some of the things he thought about; absolutely he kibitzed much of the stuff I did. Together we got it done, but he was clearly the senior person on the project. Then we added another person, that was Bernie Cosell, who was the third programmer, and so originally there were three of us who were programmers - Willy, Bernie, and I.

O'NEILL: Okay.

WALDEN: And then on the hardware side Ben Barker came to help. He was one of Severo's students. Severo taught part-time at Harvard. He taught a hardware course. Ben Barker was one of Severo's students at Harvard. He was somehow transitioning from being an undergraduate to a graduate program or had been a graduate student already and was not finishing his Ph.D. yet or something or other. He did his Ph.D eventually. Ben is now Senior Vice President of the BBN parent company in the business development area. So Ben then joined the group as Severo's, I think partner is the right word; Severo was the senior engineer, Ben was his junior partner. Willy was the senior software person; I was his junior partner. I think that is the fair way to say it, as

opposed to Willy was somehow the manager and I was subordinate, because Willy is not inclined to manage anyone. He is a researcher, but he was clearly the senior researcher.

O'NEILL: Was that pretty much the group then that responded to the proposal?

WALDEN: The group that responded to the proposal was, then, Frank, Willy, Severo, me, I'm not sure if Ben was here or not by that time, Hawley Rising, who had come from MITRE and was helping manage Frank's division. I would guess Alex McKenzie was not working on the project yet. He probably started after we got the project. And then a few others from around the company who kibitzed, such as Joe Markowitz and Dan Bobrow and others who participated. I can't really remember who, Frank might remember better. But it was a relatively small group. I would say it was sort of six full time people and ten altogether working on the proposal, something like that.

O'NEILL: Did that change significantly once you got the contract?

WALDEN: Not a lot, no. Because really the software effort got augmented by Cosell. The hardware - Ben joined if he wasn't there already. I think he was, and maybe Marty Thrope came. He was another person in that picture. This guy, Bill Bartell. We used a Honeywell machine, and we got Honeywell to assign a person to us full time on the thought that on a very rushed project (I think it was eleven months from the start until the first delivery) it was important not to have to spend very much time figuring out the computer. So we asked Honeywell, as part of selling us these machines, if they would provide this guy, Bill Bartell. And I am not sure if we paid for him or if they loaned him. So we had him. And Alex joined fairly early. I don't remember when exactly. Probably a couple of others helping in various ways. It was basically a fairly small project. It was surely never more than, in the first year, not more than a dozen.

O'NEILL: Was Honeywell the clear choice? Was there a lot of debate about it?

WALDEN: They were a relatively, let's call it, straightforward choice as opposed to a clear choice. I mean that in the following sense. It wasn't that they were the only possible group. We took a look at all of the minicomputers of the time, the real-time-oriented mini-computers. (Come to think of it, I believe Alex helped with this so he must have been on the project early on.) We looked at how their interrupt system worked and how many data channels they had. We talked to them about whether they would help us build boards to slide into their machine, because we had to build special interfaces to interface to the ARPANET protocols to the modems and the ARPANET protocols to the host computers. We looked at everything that seemed sensible. We looked at the DEC machine, we looked at the Honeywell machine; I suspect we looked at some kind of an IBM machine. I would say that it was our conclusion that the Honeywell was the best machine for that job that year. Could you have used some other machine? Probably. Was the Honeywell the best machine? Would it have been the best machine the year before or the year after? Who knows. Probably not. So it was a straightforward decision. There wasn't a lot of debate about it, but it wasn't somehow that "it was the only good machine and the others were no good." It was just the winner of an evaluation.

O'NEILL: When you were first starting out, were you aware of the work of Paul Baran at Rand and others at Rand, and at NPL and various other things that were going on?

WALDEN: Certainly, Bob Kahn would know better. But I suspect that Bob Kahn had read all of Baran's stuff. He was in the community. It was known. I can't remember the first day I heard Paul Baran's name. But I'd be surprised if I didn't know it in those days. Because some of the issues that came up were issues like: "Do we have to do, worry about some of the stuff Paul Baran is worrying about. What is the RFP going to say? Are we really concerned about reliability in the face of nuclear attack?", and that kind of stuff. The NPL work (there had been a conference somewhere - one of the operating system conferences a year or so earlier [Gatlinberg] - yes, the Gatlinberg conference may be that Roberts and Donald Davies were both at) - so that work was known of.

O'NEILL: Were you at that conference?

WALDEN: No. But I think it didn't influence us very much, any of that. At least it didn't influence me. I knew of it but didn't know that much about it. We were building a very pragmatic computer system. How do you, with this computer, this programming language and this RFP, build that system? My view is that packet switching at the time was something ripe to happen. There was the earlier work of Baran, the computers had now gotten fast enough and had good enough memories and were cheap enough so you could actually implement in software algorithms, which ran fast enough to get the job done. There were attempts to build things at NPL. Something was going to happen. There had been the experiment from Lincoln to SDC that Tom Marill and Larry Roberts were involved in. Bill Mann, by the way, who eventually worked here also worked on that project and there was another person. There were two people. When I was at Lincoln Labs, in fact, that project was going on, and Bill Mann and Will Crowther and the other guy on the project, whose name I can't remember right now, from Tom Marill's company, Computer Corporation of America, were rock-climbing buddies. And so when they were around Lincoln working on the Lincoln end of the Lincoln SDC link, Willy and I would talk to them in the hall. So we sort of knew about that earlier experiment.

O'NEILL: Did you ever use the system at all?

WALDEN: Somebody used it. It was on the TX2, and the TX2 was in a different division. It was down in the basement and I was up on the second floor, and so it didn't have a lot to do with us, but we knew about that work. I can't say that it affected us too much. I think the way these things happen is that the world is ripe for a change; it is ripe because the technology has come to a certain point. It is ripe because people believe it should be done, and then somebody goes and implements it. It becomes an engineering problem as opposed to a theory problem. My view is that mainly what we were doing was we were very pragmatically doing engineering. We had to send these bits down the wire: how do you put a header on the front; how do you put a trailer on the back. There was theory of how you put error correcting codes on it. Bob Kahn knew that theory

and told us what it was. There were some constraints: this is the way that the 303 (or the 301 or whatever the Bell modem is) has to be interfaced to, but after that it was all pretty pragmatic. Not lots of theory of coming from someplace else.

O'NEILL: OK. These other works that went on before - did you use them as more of a theory?

WALDEN: I didn't think about them. Probably Bob knew about them and either accepted them or rejected them, I suspect, or led us to think about those parts of them that were relevant. I was writing and designing software. We had a fairly free hand. Very early on one of the first reports we were supposed to produce was what eventually became BBN Report 1822, which was the original specification of how a host communicates with a packet switch. We wrote it. We wrote it partly keeping in mind what we had been told the hosts wanted and quite a lot in keeping in mind what was going to be possible to implement and what made sense to us. We then took it to the committee of the hosts, the host protocol people, the host representatives, and they kibitzed it a lot and told us where they didn't think it worked and there were debates on that. Some places they said they didn't think it was going to work and we said we thought it would work. And then they were later right.

O'NEILL: Did anyone in your group have specific communications background?

WALDEN: Well, Bob Kahn. Let's look at it in a couple of ways. First, in our group Bob Kahn had a Ph.D. in Communications Theory, I think, from MIT. His Ph.D. is in electrical engineering and math-ish stuff, but he came in the era when studying communications theory and stuff that Shannon and all those people did was what he studied. So Bob absolutely had an academic background in communications. He had worked at Bell Labs, in fact. Frank and Will came from Lincoln Lab where they had been building communication systems. They were building the systems which had pointed radar antennas at satellites, and at the moon, and at various things, to do some of the early satellite communications experiments. So there was a background in communications. I think, however, that the fundamental background that I brought, that Frank brought, that Will brought, that

Severo brought, was not communications knowledge so much as knowledge in building practical, real-time systems. The fact that it happened to be a communications system is a little bit incidental to the fact that this was a group. Frank was a person with vast experience in building successful real-time systems, systems that did whatever they were supposed to do, second-by-second. ~~Frank's design philosophy has always been a defensive design philosophy. Assume the worst. Not to a point of silliness, where it doesn't work because it's got too many cycles busy assuming the worst. But don't assume the best. Let's put it that way. He believes that unless~~ you are a defensive systems designer, there'll be things you'll have a lot of trouble with. And the other thing, of course, that Will and Frank and Severo brought is their native wit; they are smart guys. There wasn't much theory in how you build a packet switching network. There was a communications theory, but that was all pretty abstract. One just got out there and did it. All the stuff that is now taught in courses in communications about networks and protocols and all of that, I would say we were mainly (as part of the entire community of the host people) inventing it. The academic analysis tended to come later in a lot of cases: how one does acknowledgements and how one does timeouts. A lot of it, the early implementation, was not so great. All of the original protocols eventually got replaced. They were useful in their time, but as time went on one understood what the shortcomings were and some of those could be upgraded and some of them had to be replaced.

O'NEILL: As the work got going in the first year especially, how did the group work together? Were there weekly meetings or was the group small enough to keep informal? Did the software people go off separately from the hardware people?

WALDEN: It was too small. We all sat in offices that were right next to each other basically and we worked together constantly. There were often meetings. I would be surprised - I don't remember such a thing as we had a weekly progress meeting. We probably were more in tune with the progress than that; we probably did it hourly. There were certainly meetings. "Let's get together and think about the Honeywell machine," or some subset of us had to get together because our original set of ideas of how you ran the modem interface didn't

seem to be working and that had to be fixed. But it was a very small group working together all the time. We were working lots of hours. There was, as I remember this started quite early, a series of informal working notes which I think we called "the impguys notes" and we used that in addition to all the informal discussions - people would write up "here's my thought on how to implement this, here is my thought on this algorithm, here is my thought..." and so on.

TAPE 1/SIDE 2

WALDEN: Probably twice as many people got the notes as actually were working on the project because everybody was interested in what was going on. And we used that as part of our communication mechanism. As time went by, as other people got added to the project, as we tried to get more orderly, as we began having to deal with (when the network was actually going in) well, this host person wants that - we would make a list of what needed to be done and so on. But we used the informal note series quite a lot, I think, and as I remember it started fairly early on. I may be wrong, it may have started a little bit later.

Certainly Severo ran the hardware effort with Ben helping him. Severo was kind of the assistant project leader in some sense, this is my memory, although Bob was a very, very influential person on the whole project. I don't know what Severo was doing with Ben all the time and eventually Marty, maybe they were having meetings. On the software side it was Will and Bernie and me. Will is not inclined to be a manager. And we just worked together, we worked together all the time. Constantly. We were in and out of each other's offices and helping each other debug and everything. Probably Frank asked us once in a while how we were doing. I suspect he did, I just don't remember it as really that relevant. It was so obvious. We'd run in and say, "Look, I got this running. Somebody come and type on the teletype." "This is exciting. Something is cycling."

O'NEILL: Yes, that is one of the nice things about a small group.

WALDEN: And so it was pretty obvious that progress was being made.

O'NEILL: Can we shift a little bit to talking about the relationship with DARPA and the IPT office?

WALDEN: Are you going to come back to this? I mean, as the project went on then things changed eventually. I would be glad to talk about that, too, and how it got managed.

O'NEILL: Why don't we continue on with that?

WALDEN: Well, certainly in the first year and then people got added...you should ask Frank...you really do need to interview Frank. He is the one who was running the project. Eventually we got it all implemented. Then Ben and I went out to California with the first packet switch - Ben as the junior hardware guy and me as the junior software guy went out to California with the first packet switch. At UCLA, from their side were Steve Crocker, Vint Cerf who were their host protocol team, a guy named Mike Wingfield, who built their hardware interface. Mike is at MITRE Bedford, I think, now. A couple of months later, there was a big issue of network performance. Bob and I went out and spent several weeks in California running network tests. There was a paper that resulted from that. It was quite a famous paper in the early days in the ARPANET. I can probably find it. I think he and Crowther might have written it. As we eventually put the packet switches into the other original sites - into SRI, UC Santa Barbara, and Utah - I was the one who did the software releases, at least some of the time, as the junior programmer. I took my paper tapes on an airplane and would fly to Utah, go to the place, put it in the machine, start it running, fly to SRI, put in the machine, start it running, fly to Santa Barbara and fly on to UCLA. And then I was also the one who maintained the software again as the junior software person. Well, I am sure, off thinking about additional things and I was probably thinking about additional things as well. But I remember that up until September of '70, I had the program listing beside my

telephone at home and would get telephone calls at home whenever something stopped. My telephone number was quite literally on the front of the original packet switches and they called my house over here in Allston and I would talk to whoever was there, using my listings. So that's what happened. Then, I went away to Norway for a year. I lived in Norway, where I incidentally helped build a packet switching network. The second packet switching network was built for the Norwegian Air Force by NorskData. And I was the project leader for that, based a lot on some of my experience here. (That was '72?). September '70 to September '71 I was in Norway. September of '71 I came back to BBN. I had just quit and went to Norway. I wanted to live in Europe. After a year I decided I wanted to live in the US again, so I came back to BBN, rejoined the project, and gradually over not too long - now the project was getting a little bit big. By the time I got back Bob was project manager in the other division and was probably by then already having discussions with Roberts about going to DARPA. This is all a little bit of speculation on my part. Somebody needed to work on the software and be really responsible for it and that is not Willy's inclination so somewhere in there, probably within a few months certainly no more than a year after I got back, Frank suggested I should manage the software effort. It wasn't a much bigger team, it was Crowther, Cosell and me, a couple of others maybe coming eventually. Crowther, the relationship between Crowther and me, having been what it was, that is, he was my mentor basically when I got out of college and I had never managed him, one couldn't think of him as a subordinate certainly. He was much more experienced, much smarter than I was. But we got along fine. And so I began getting involved in "Well, we need a document for this and we need to have a schedule for that," and evolved eventually into running the software effort, hiring additional people. Then when the network management part began to get more formal, I became responsible for that because network management was a lot of software. Alex McKenzie joined me somewhere in there (he had been on the project for some time) as kind of the assistant network manager. That went on for several years with a variety of projects, the original satellite stuff, the original terminal IMP stuff (although, I think the Terminal IMP project started while I was in Norway). Frank remained the Principal Investigator. Probably sometime in there evolved that the hardware projects almost reported to me, and as time went by, I sort of, because that was my inclination, took over more and more of the management of the project day by day. And eventually I was defacto the manager of the ARPANET effort at BBN day to day, month by

month, and year by year, on what should we be doing. Frank had a division to run - Frank, of course, had a big influence. BBN being the kind of place it is, the whole staff has a big influence anyway - there are a lot of votes on what we should do next. So, that is the answer to that.

O'NEILL: You mentioned actually going to these sites and putting in tapes. What kind of support did you get from the sites themselves. Were there people there to help? You mentioned UCLA.

WALDEN: Each site typically had a professor of some sort who was responsible for this effort, who had the DARPA contract at that site. Typically there was some kind of a hardware person and some kind of a software person. So the hardware person had built the host-interface-to-IMP hardware interface, and the software person had built the original software interface to the IMP, including maybe the host protocol, and there were those people at each site. And they were typically students maybe or staff, like a University has. They were all pretty easy to get along with.

O'NEILL: Were these first four sites chosen by you? Did they volunteer?

WALDEN: They were chosen by DARPA. We were the fifth site as I remember. Because it became obvious that being connected in would be useful.

O'NEILL: Yes, I bet!

WALDEN: I think the right way to view what we did was that day to day DARPA was the manager of ARPANET and we were the operator of the ARPANET. That is probably the right distinction. For many years, BBN was the operator of the Defense Data Network and DCA was the manager of the Defense Data Network. People would call us up and ask how you get on the ARPANET and we would refer them to DARPA, or we

would refer them to whatever government agency DARPA was using to help them. And then they would tell us OK talk to them and we would send them a document and explain to them how to connect to the net.

O'NEILL: Did you personally have a lot of contact with people at DARPA?

WALDEN: Not originally. I think there are several different things to say. One is that in the early days when we were actually building the first packet switch we were all working lots and lots of hours here and there were some review meetings. They came to visit us - where they is DARPA and as I remember it was members of the host community (people like Rulifson from SRI, Crocker from UCLA, and whoever was at Utah) - so these people would come and have meetings here and we maybe went to a meeting someplace else once or twice. I suspect there were meetings that Bob Kahn or Frank Heart went to that I didn't go to. I was busy implementing. So I would say that in a very natural way that as if there was a meeting at BBN or if there was a technical meeting somewhere else the DARPA people would be there and we would be there. So in that sense, there would be communication with DARPA. As far as how the program was being monitored on a contractual basis, I don't really know what was happening the first year. Probably Frank was sending some kind of a monthly report off to - in fact, there was, I know there was, of course, it is all coming back to me. There was a quarterly technical report, which probably Bob Kahn originally wrote. Those are all probably on file somewhere in wherever the place is you send government reports to. And I certainly contributed to some of those quarterly technical reports. Because I wrote things up probably a little more than Will wrote things up. So I would write up software pieces and eventually I became responsible for writing the Quarterly Technical Reports as I remember. And then after than Alex became responsible. It is the kind of things that went to the writers. So Bob (Bob's a writer), so he was certainly the first person that was the internal person that wrote them and whether Frank reviewed them and how they finally got finalized, I don't know. But as I remember Bob wrote the early quarterly technical reports. I don't know what kind of a quarterly or monthly financial report there was. There is a guy at DARPA, Al Blue, by the way if he's still around. Do you know the name?

O'NEILL: Yes, he has been interviewed.

WALDEN: He would remember how that worked in those early days because Al Blue was a DARPA guy monitoring the contracts.

O'NEILL: OK. So he was monitoring the contract with you.

WALDEN: In the very early days, there was Larry Roberts, there was Barry Wesler and there was Al Blue at DARPA. There wasn't anybody else. And Roberts and Wesler weren't inclined to anything administrative. Al was office administrator. He is still living down in a mountain in Pennsylvania somewhere?

O'NEILL: I don't know. I wasn't the one who interviewed him, but I know that he was interviewed for the project. How did ARPA let you know what they wanted? Were they calling these meetings, these review meetings?

WALDEN: I think you should read this. I haven't read it recently, I put my memories in here. But you should read this because I think it talks some about these meetings. As I remember, there was a spec, the RQP asked for some certain things. We proposed our answers to those and DARPA reviewed them in some way. And as I remember it, oftentimes it was more than just DARPA, it was the community, it was the host people. They would say, "We don't like that." Then we would have arguments. So it was very much a collegial community debating the correct technical solutions. Then we come back and we design a little bit more and propose something and sometimes our argument would carry the day or our stubbornness would carry the day and sometimes we would get convinced of something else. Again, I can't remember without reading what I said here or going to talk with Willy. Bob would probably remember. When the decision was made, for instance, to go to multiple hosts, did that come from them or did that come from us? I remember it as a major eyeopener and it was a surprisingly easy implementation problem. We noted, quite quickly, "Well, we can just use all this code

over and over, call it a lot of times... with different parameters." I remember that as being a fun revelation, both the eyeopener that you needed to have multiple hosts and whether that was just something that became apparent to everybody or whether somebody came and told us this is now the economic thing, I can't remember. The host people began saying "But we are going to connect multiple hosts." Then the idea that it could be implemented easily was also kind of exciting. And right after that, I remember, we figured out almost simultaneously that well, if you are going to handle multiple hosts at a site, make the program do that, then all of the functions inside the packet switch can be pseudo hosts, what we called fake hosts for years, and they can just communicate like any other host. In fact, they didn't exist. We could use all the protocols to do the data collection, to do the control of the packet switches. So every host had actually had four real hosts on it originally and four fake hosts. The fake hosts were internal software hosts doing various control and monitoring functions. That seemed neat at the time. It gave us a chance to use co-routines.

O'NEILL: Do you remember any instances of disagreeing with any direction that was given? In this case it sounds like there was an agreement.

WALDEN: No. In that case, I think there was a lot of debate. I think that there was, I am sure there was disagreement about aspects of how you communicate with the host. Was the end-to-end acknowledgement system going to be sufficient? If there was not disagreement there was at least doubt. The so called RFNM (Ready For Next Message). Was that an adequate flow control mechanism - cross network control flow control mechanism - there was certainly doubt, I suspect. And if you read Kleinrock's papers and other papers, I suspect, you would get other peoples' points of views on that. And I think there were certainly cases there where we were probably stubborn. Or, just there wasn't any time to do anything different. I am not sure which it was, probably a little of each, so we implemented what seemed ok. In the end-to- end and flow control area, it almost immediately became obvious to anybody who hadn't believed it before that what we had originally implemented was insufficient. And what we had specified to the host and what they had implemented had a lot of ramifications. They were now done and now it didn't work all that well. So for the first year or two years, quite

a long time, there were rules, informal rules in the network that you didn't pump stuff into it from the hosts so fast it swamped the net. Because we all knew the algorithms didn't work, I mean they broke, patently they broke when you pumped stuff at them too hard. And so this interim solution worked in that network -this is the kind of thing that somebody who is working on proving programs correct and stuff can't imagine. It actually was quite useful in those early days of experimentation because everyone agreed not to break it. You know how to break it; let's not break it. Meanwhile, let's go figure out how to fix it. And so then there was a reimplementaion of ways in on the part of the host interface to take care of it. The tables which were now kept at IMPS at the destination to keep track of how many buffers were allocated and so on so you did not get any deadlock conditions at the destination.

O'NEILL: So it was on the order of a gentlemen's agreement. Just not to go break it.

WALDEN: Engineers' agreement. Right. Everybody knew that all you had to do to break the network was to send traffic at it too hard. So there wasn't much thrill doing that. All it did was make it unpleasant for everybody else. So people didn't do that. But I think this is probably a technical area, again you should talk to Kleinrock and Kahn, but this is a technical area where surely what we implemented - I am virtually certain there was some doubt. It soon became apparent it was not right. And I would guess that Bob Kahn might say that he understood all along that what we were doing was inadequate and that Crowther and I and the programmers were just too stubborn. But I don't know if he would say that or not. You could ask him.

O'NEILL: When you talked about the collegial atmosphere, are you talking about dealing with the hosts or does that include the other contractors?

WALDEN: What other contractors?

O'NEILL: Well, like the Network Analysis Corporation. Would you distinguish between those two communities?

WALDEN: Well, what I would say is that the groups involved were Network Analysis, Network Information Center at SRI, the whole committee of hosts, the host protocol group led by Crocker and Cerf as I remember, certainly Crocker quite a lot. Us. The telephone company. DECCO. I think DECCO was how those telephone lines were being supplied in the early days. So DECCO was the group contracting for the phone lines, but we were probably dealing directly with the telephone company. And Bob Kahn probably did most of that dealing for us with the telephone company. I would say that the collegial atmosphere didn't extend to DECCO because that is a fairly traditional procurement of Scott Air Force Base in Illinois as I remember. I would guess they were business-like but certainly not collegial. The telephone company, I was not in close contact with them and I suspect that they did a decent job. There was certainly some guy in the telephone company that Kahn dealt with that was maybe their account manager who was trying to be helpful, but I think we felt like the telephone company was a pain in the butt a lot of the time. It took a long time for the lines to get in and all that. Then there was Kleinrock, the network, whatever he called himself.

O'NEILL: Network Measurement Center.

WALDEN: Network Measurement Center, right. Well, I would say candidly that with the host people and with DARPA it was all very collegial. Certainly, everybody came to all the meetings, although we probably thought what Kleinrock was doing was relatively academic. The interface with Network Analysis was much more between Larry Roberts and Howie Frank in Network Analysis than it was to us. That is, maybe Kahn talked to them but I don't remember that a lot. What I think happened was DARPA was getting the network requirements. He was asking Network Analysis to do the network analysis and then that was being brought to us. I don't remember there being so much discussion there. Being computer programmers and hardware designers as opposed to academics or opposed to theoreticians, let's put it that way. We probably viewed what they did as

somewhat academic, too. The Network Information Center - I always kind of viewed trying to grab off the bigger and bigger pieces of network work. This is my candid view of it. But in a meeting everybody was well behaved. Civil and undoubtedly performing useful function and our point of view a little bit parochial. But with DARPA and with the host protocol people I remember it being very collegial, not always calm, but like scientists arguing with each other about what is the right answer.

O'NEILL: Did Roberts or ARPA, IPTO whatever, did they arbitrate these disagreements?

WALDEN: Probably, I don't remember exactly. If it needed to be arbitrated. Probably a lot of it just got decided. We figured it out. How does design get done anyway. It is a question which comes up often. We work with big companies. We "team" on some kind of project with IBM or Hughes or something. The question that we spend all this time working on when you are writing the contract is who's the boss? and who's not? But that's not how it actually works: mostly what happens is you sit in a room and argue about it until you all agree about what the right answer is. And that is mostly how I remember it happening. We were in there (I was, certainly others were) busy writing host protocol RFCs, putting in our two cents' worth on how host protocol design should be done. Having a fairly major influence, I suspect, because we were as close to it as anybody. We put a lot of hours of thought into it. So I think of it more like a standards committee but not that formal. How do you arrive at these things? Standards committees have the disadvantage usually that they are after the fact and there are lots of rival points of view; and so everybody's trying to keep a little bit of his point of view. This, perhaps, is a somewhat cynical view of standards committees, but there is probably some truth in it. There was nothing here before so it was a bunch of people trying to find a standard without anything having gone before. And there were real practicalities, as I suggested. You couldn't just tell everybody, "Oops, we made a mistake - go reimplement it all tomorrow." They didn't have any money to reimplement tomorrow. So one lived with those mistakes. The whole transition from the original set of host protocols to TCP/IP was a very long transition. Because everybody thought they were done.

O'NEILL: Can you characterize how the ARPA office influenced the development?

WALDEN: I'll give you another point of view. And I am not sure if everybody would agree with this or not. I believe that the ARPANET project as a whole went very rapidly and relatively smoothly despite all the mistakes we built in because there was a single contractor who was responsible for the backbone net and that was us. We had a contract. We couldn't change our mind too much. We couldn't have our mind changed too much because there was no time - we'd overrun. I would say, but then I don't know if anyone would agree with this, that the group here because of its role as the prime contractor to the backbone net was clearly the most influential group. And plus it had some pretty strong personalities with Heart and Kahn, and so on. I think that, you know, if Roberts gave us direction, we would argue with him and then do what he said eventually. This is a science community, this is not a group that is inclined to be "yes" men, ever. I think mainly what he tried to do was try to make sure we got done on time on, didn't get too distracted, didn't get too many conflicting signals. But did listen to what was important to the other people. When the "gang" (and I use the word in quotations) of host protocol people would arrive as a big group with their concerns, and shout at you for a day - that was pretty influential.

O'NEILL: Getting shouted at usually is, yes.

WALDEN: So you go do something sensible.

O'NEILL: Were you encouraged to publish and try to get even more information out into the community?

WALDEN: I don't know if we were encouraged to publish or not. We did. A lot. I think publishing is the kind of thing that depends on the individuals. If the people you happen to have happen to be people who like to write, you get a lot of publications. Heart, Severo, Kahn like to write. I discovered then that I like to write. Crowther got taken along as co-author often. He actually wrote but he would never write just for the sake of publishing. Plus, because of Roberts or whatever, DARPA, I don't know, this was not a place where the military was busy saying publication is bad. Plus, there were these note series. There were the RFCs.

TAPE 2/SIDE 1

WALDEN: We were talking about writing. It occurs to me that Roberts wrote. That probably was a big influence. Roberts is a writer, he has always written. Wessler, Kahn, so I think it was just a part of the culture of this collection of people that one wrote.

O'NEILL: When you talked about the host protocol group was that the same as the Network Working Group?

WALDEN: Yes, exactly the same. You can put Network Working Group.

O'NEILL: When you gave, I don't know that you personally gave presentations, but in response to your articles or to conference...

WALDEN: I personally gave them.

O'NEILL: OK. Was there a large level of interest outside of the host community?

WALDEN: I think very relatively early. Getting invitations to come and speak to this group. Go to visit somebody at IBM. Go give a presentation someplace. Give a presentation in Europe. You could look at the bibliography, but quite early on there were presentations being given. And, of course, there was the effort, it must have been in 1972, the first five papers were in the Joint Computer Conference.

O'NEILL: It was 1970.

WALDEN: Seventy. So very early in 1970 and then again in 1972. It was the effort to put them at the Joint Computer Conference. Well, in fact, if you remember, the original, probably both sets, but certainly the original

ones in 1970, DARPA bound in a government cover distributed to the world, so clearly Roberts wanted to see it out there.

O'NEILL: That's what I was wondering, if they were encouraging either to write it from the very beginning or were they just taking advantage of the fact that they were there.

WALDEN: You had better ask them. It seems to me though looking back that they clearly were trying to proselytize for packet switching. I suspect one of the things we were trying to do is find users. But for Roberts to bind those up and distribute them, make them available; it was clearly decided, I mean I don't know if it was a plan or it was just so clear to everybody that they knew it. It was decided that Roberts was going to have a session, there were going to be these five papers in 1970. There was going to be one from each of the groups. There was Roberts and Wessler, there was our paper, there was the Kleinrock paper, there was a NAC paper, I can't remember who the fifth one was.

O'NEILL: There was one by Cerf, Crocker, and Carr.

WALDEN: Sure, the working group paper. Exactly. So there you have the five key areas.

O'NEILL: When I think about someone proselytizing for something...

WALDEN: Cerf, Crocker, and Carr. Steve Carr, right.

O'NEILL: Yes, I think of having some resistance or people needing to be persuaded.

WALDEN: Absolutely. Absolutely. The packet switching is clearly a technological change. Message switching and circuit switching were what was before. The military had used message switching. Telephone companies

were all for circuit switching. And that packet switching could even work was doubtful a little bit to some people. And there was great resistance to any change. The message switches did things like know about special codes, addressees, and all that specialized stuff.

O'NEILL: You're thinking of the AUTODIN?

WALDEN: AUTODIN lines, yes, all of that. What came before packet switching, then all the airlines that the government used was message switching. Where you sent big files and there was a lot of application specific addressing in it. Yes, there was resistance to that.

O'NEILL: How did you see the resistance? Can you give me an example of it?

WALDEN: I think there are two kinds. One kind was probably in the military with the various services saying "we are going to do our own thing." The kind that was probably more influential, that we cared about, was the academic disbelief. It is a thing that we have today. We propose: this is the right way to build the parallel processing computer. And one group of people gets up and says -- Amdahl gets up and makes a speech about how you can never go faster than linear or something (I don't know - whatever Amdahl's fake law is). So you have a lot of people propounding laws. You can never take advantage of it. Then there is a different set of people saying things like "the hyper cube architecture is the only one that can ever be made to work." And here we are busy implementing this one, we can see it works, and everybody is telling us this can't work and it is not useful. So I think it's more...the thing that I feel we probably cared more about was the academic doubt. The rival theories of how one should do communication.

O'NEILL: Do you remember who was putting forth the rival communication theories?

WALDEN: The academic community a lot. I am serious. I think the academic world. There are other scientists and researchers out there building packets switches, building circuit switches for telephone companies, people at Bell Labs, people at IBM saying "Who needs packet switching," and what is packet switching going to do, right. Packet switching is going to facilitate communication between different kinds of computers in different companies. That can't be good. I am not saying that the person at Watson Research Center who says that packet switching is not a good idea is only touting the company line, but he is working in a department that is working on host to host communications using SNA or whatever the predecessor of that was, some telecommunications protocol. Time division multiple access or something, I don't know. So, naturally, that person is going to be explaining how the only way to really do it is with 3750 or whatever, those 2780 front ends. I don't know what those numbers are. It is that kind of thing. Just technical rivalry. This is the sense I have.

O'NEILL: OK. Did these come out at these conferences? Were there debates or discussions?

WALDEN: I probably never did go to any conferences so I don't know. I always felt we were defending ourselves.

O'NEILL: OK. I guess I haven't really seen a lot of articles that explicitly say that packet switching is a bad idea. So I was trying to get a feel for what the arguments were and where you were hearing them. Whether it was being written or just discussed.

WALDEN: It's the feeling I get. The vendors, certainly the vendors were saying this is not a great idea.

O'NEILL: Even when you were including their systems as hosts?

microsecond and millisecond times you had to do your switches with hardware. Well, switching with hardware can't be very complicated. It's that simple. Suddenly, we could do our switching with software. I always knew, I think we always knew, although we didn't always say this out loud, that the part of packet switching which is buffering the individual packets was an engineering tradeoff that was right for the time. The fact that we were doing switching with software was the bigger key. And what we are moving into now with the software switching, without the local buffering, things go flying through external buffers and there's fiber cables switched across the country on microsecond bases, this whole gigabit networks technology. I think that from very early on we understood, probably Kahn understood it before that, I understood it, let's put it that way, from very early on that the key thing here was that we were doing the switching with software. Because the computers now went fast enough. As it happened, the memories were cheap enough so you could do the local buffering, but it was only a matter of years - whether it was ten or twenty - before the hardware would become cheap enough so you wouldn't necessarily have to buffer the stuff locally. But you still had the software based switching. Computers got to that point, you could do the switching with software. Didn't have to take it in and store it for an hour and look it up in some big table somewhere. Well, the line costs were coming down. So the fact was that you were using them a little inefficiently. The bandwidths were going up and the relative costs were going down so the fact that you only used them at 60% of capacity wasn't a disaster. Message switching was optimized for let's fill the lines up, every hour of every day. I think there is another thing that I believe Roberts understood, Kahn, and those people who are real thinkers understood early on, certainly the rest of us realized quite quickly, which is that putting in this infrastructure would change the way the world worked, not the communications world, but the way people worked. From the first time we sent a message across the network or wrote a paper across the network, none of us had any doubt that what you are seeing today with thousands of distribution lists and virtual networks and worldwide queries. It was going to happen. So we thought we were changing the world, I think.

O'NEILL: Did that help in putting in the long hours?

WALDEN: Yes. It was fun to implement. I think it was a thing whose time was ripe. The government did a good job in selecting this group because it was one that would make good technical tradeoffs rather than being dominated by selling how you maximize the sale of Raytheon or IBM equipment. It was a good tradeoff of practical and academic, but some other group could have done it too. It is quite clear.

O'NEILL: I was going to ask you more about the Completion Report and any follow-up you might have to it, but you have already mentioned that you have written all that down.

WALDEN: You can ask me about it.

O'NEILL: Did the Completion Report fairly well reflect your views in '78, '79?

WALDEN: It reflected our views and DARPA's view, I'd say. Let me give you a little bit of history of the Completion Report. Somewhere in here when I was managing the ARPANET for BBN, Craig Fields had gone to DARPA when Licklider went for his second time. After Roberts left, Licklider came and Craig came with him. Craig probably had a Ph.D. in psychology, Licklider is a psychologist so Craig came. Craig, after a little bit, was made ARPANET manager from DARPA's point of view, IPTO, Information Processing Technique Office's point of view. So he and I used to talk to each other on the phone a lot. I was managing a project here and he was managing there. And, in fact, the transition from DARPA to DCA was done on Craig's watch at DARPA. Somewhere in there, he got the idea that there should be an ARPANET book. And as I remember it, I think I remember correctly, he called me up and said we should do this. Anyway, either he then suggested to a lot of other people that they do stuff or I suggested it. Anyway we had probably at least one meeting, maybe here, of people from the various hosts, and a number of people drafted chapters or, I guess they didn't draft chapters, they were given assignments to go write chapters. That's what happened. Then nobody did anything.

the rest of the world. Probably in the intervening years, I have come to believe more that they were actually right. That we were more wrong than we thought we were in the early days about some things.

O'NEILL: Specific technical things?

WALDEN: Specific technical decisions. In general, I think that the global architecture of it is pretty neat. And I don't think one has to apologize much for not getting it right the first time and getting it good enough to be used, and all of that. But there were a lot of ideas on the outside that we either didn't accept as readily as we might either because we just didn't have the time or because we wanted to be right.

O'NEILL: Did BBN have involvement later with the Internet and the switch over and all that?

WALDEN: That's all in here, that's what this is about, how the evolution of the Internet goes. And Alex is better to talk to about that. I really wasn't involved with the network from 1978 on, from 1978 to 1982 I was kind of out of touch. I was doing other things at BBN, and then when I came back, so that that part was in my part in BBN, other people were doing it, Steve Blumenthal and Alex McKenzie. Alex is a good person to talk to. He is here, (at BBN) he is downstairs somewhere, in this building. He is here and he's been involved almost as long as anybody. In fact, Alex probably in the first decade is the number one player in the whole area of network management. Alex came on the project when network management was originally being implemented. He was involved in network management. He was involved in the International Standards Committee's taking the ARPANET protocols and making them match some of the European protocols, for instance. And all the network management stuff still reports to him. Alex is an Assistant Division Director in BBN STC who is responsible for the communications research. The continuation, in some sense, of our original ARPANET contract - it is not our literal continuation - but we started working in communications for DARPA in 1968. We are still working in communications for DARPA doing continuing research - that research reports to Alex. We

are proposing work in gigabit network, we are working in replacing the ARPANET with the Defense Research Internet. The satellite networks that still go to Europe are in Alex's area.

O'NEILL: OK. Did BBN have any communication projects before you responded to the RFP?

WALDEN: I doubt it. Bob undoubtedly had some kind of funding from DARPA to help them think about communications.

O'NEILL: This is Bob Kahn?

WALDEN: Bob Kahn. And I think he probably came here thinking there was an interesting match of technology and technologists here. Communications is going to be important and he would put us into the communications business. Which he did. That's my guess. As I said, the key thing we had was people with real-time computer system experience. Smart people. This was the early days. There hadn't been a lot of real-time computer systems already, but Frank had done several successful ones by that time. Frank and Will came from the Whirlwind Project at MIT.

O'NEILL: OK. I didn't realize that they both been actually working on the Whirlwind.

WALDEN: Yes, well Frank worked on the Whirlwind Project as a graduate student or a staff member. Will did his bachelor's thesis on the Whirlwind. Then they went to Lincoln and joined the group that had put in the SAGE system. Then they built the Westford and the Haystack antennae pointing systems for doing radio astronomy. And then they built the Lincoln Experimental Terminal System. They were in the division that threw up the "needles belt" and bounced stuff off the little specks of metal up in the sky. So they were very experienced real-time computer systems builders.

O'NEILL: That pretty much covers the set things I had wanted to ask. Is there anything else you wanted to add?

WALDEN: Whatever you want to know is what I want to tell you.

O'NEILL: Okay. Thank you.

END OF INTERVIEW